Lecture 4
Submanifold reconstruction

Jean-Daniel Boissonnat

Winter School on Computational Geometry and Topology
University of Nice Sophia Antipolis
January 23-27, 2017

Outline

1. The reconstruction problem
2. Distance functions and homotopy reconstruction
3. Delaunay-type simplicial complexes and homeomorphic submanifold reconstruction
4. Mesh generation of surfaces

1. The reconstruction problem

2. Distance functions and homotopy reconstruction

3. Delaunay-type simplicial complexes and homeomorphic submanifold reconstruction

4. Mesh generation of surfaces

Reconstructing surfaces from point clouds

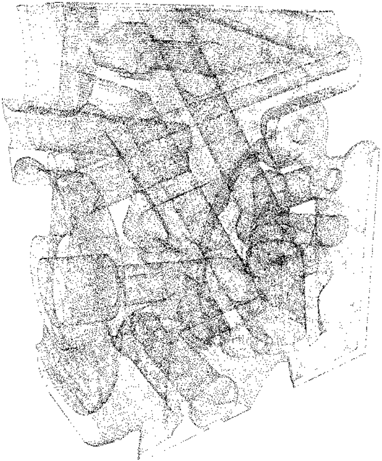

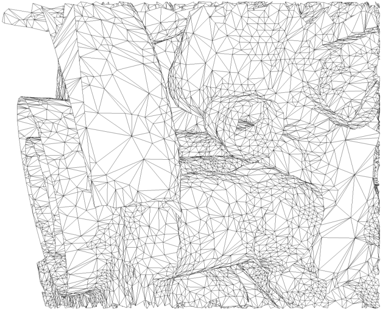

One can reconstruct a surface from 10^6 points within 1mn [CGAL]

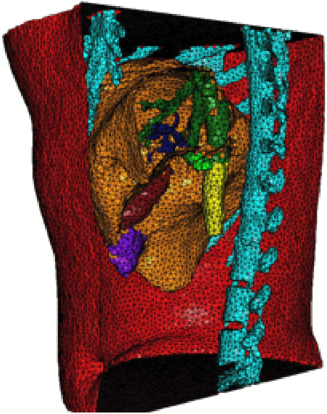

CGAL-mesh, Titane

GeometryFactory, Acute3D

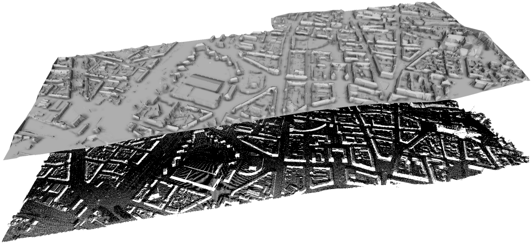

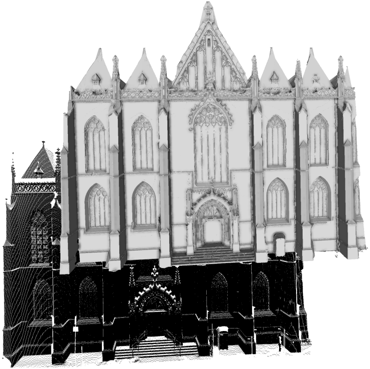

Geometric data analysis

Images, text, speech, neural signals, GPS traces,...

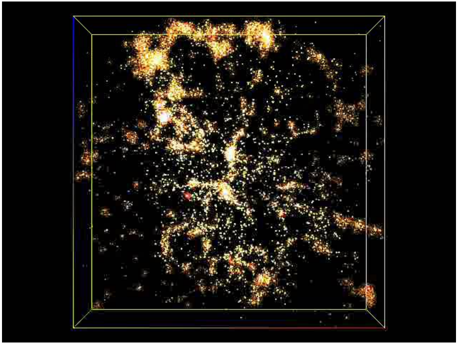

Geometrisation : Data = points + distances between points

Hypothesis : Data lie close to a structure of “small” intrinsic dimension

Problem : Infer the structure from the data

Dimensionality reduction

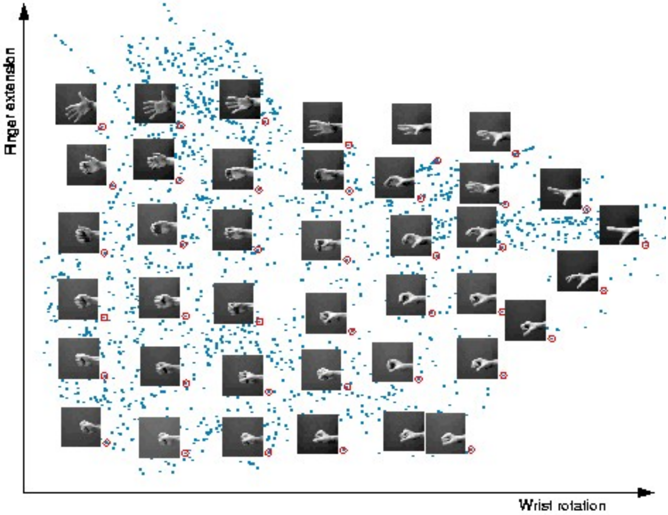

Conformation spaces of molecules e.g. C_8H_{16}

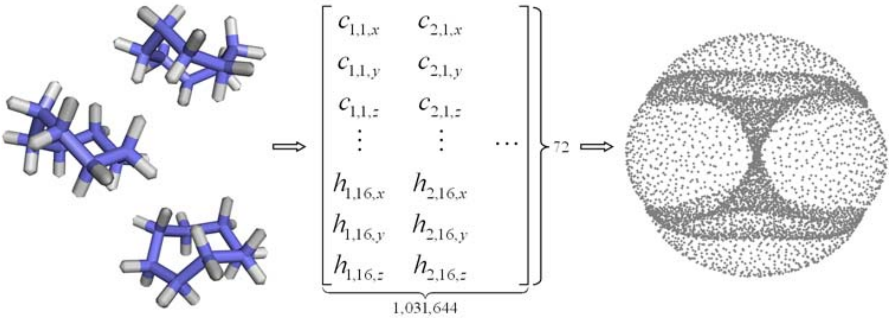

- Each conformation is represented as a point in $\mathbb{R}^{72}$ ($\mathbb{R}^{24}$ when neglecting the H atoms)
- The intrinsic dimension of the conformation space is 2
- The geometry of C_8H_{16} is highly nonlinear

Image manifolds

An image with 10 million pixels

$\rightarrow$ a point in a space of 10 million dimensions!

camera : 3 dof
light : 2 dof

The image-points lie close to a structure of intrinsic dimension 5 embedded in this huge ambient space

Motion capture

Typically $N = 100$, $D = 100^3$, $d \leq 15$

Sampling and distance functions [Niyogi et al.], [Chazal et al.]

Distance to a compact K : $d_K : x \to \inf_{p \in K} \|x - p\|$

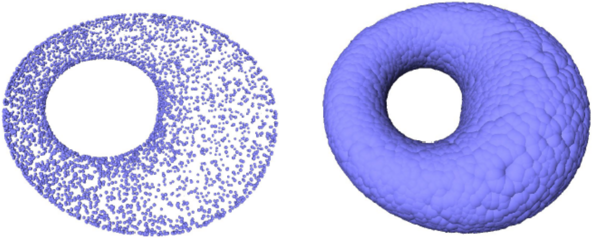

Stability

If the data points C are ε-close (Hausdorff) to the geometric structure K, the topology and the geometry of the offsets $K_r = d_K^{-1}([0, r])$ and $C_r = d_C^{-1}([0, r])$ are close for $r \in [\Omega(\varepsilon), \mathrm{Reach}(K) - \Omega(\varepsilon)]$

Local feature size and nets

[Amenta & Bern 98]

The medial axis of K is the set of points of the complement of K with at least two closest points on K

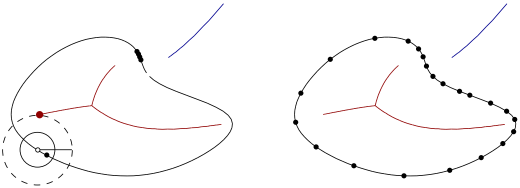

A finite point set $P \subseteq K$ is an ε-net of K if

1. Covering: $\forall x \in K,\, d(x, P) \leq \varepsilon \mathrm{lfs}(x)$
2. Packing: $\forall p, q \in P,\, \|p - q\| \geq \eta_0 \varepsilon \max(\mathrm{lfs}(p), \mathrm{lfs}(q))$ for some cst η_0

lfs denotes the distance from x to the medial axis of $\mathbb{M}$

Properties of lfs

- lfs is 1-Lipschitz : $|\mathrm{lfs}(x) - \mathrm{lfs}(y)| \leq \|x - y\|$
- $\mathrm{lfs}(x)$ is small where the curvature is large and where the thickness of $\mathbb{M}$ is small
- $\mathrm{lfs} > 0$ if S is $C^{1,1}$
  i.e. normals exist everywhere and the normal field is Lipschitz

$\inf_{x\in\mathbb{M}} \mathrm{lfs}(x)$ is called the reach of $\mathbb{M}$

Distance functions and simplicial complexes

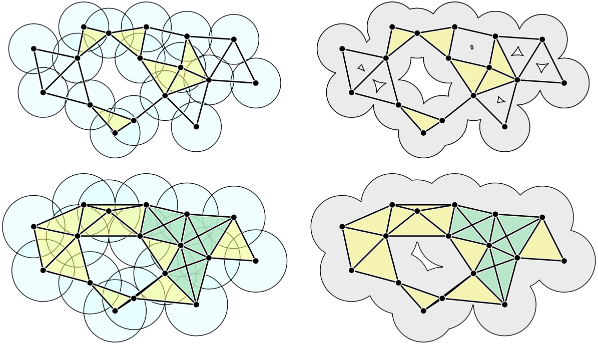

Nerve theorem (Leray)

The nerve of the balls (Cech complex) and the union of balls have the same homotopy type

Some remarks and questions

From continuous to discrete and back:

$\boxed{\text{Shape}} \rightarrow \boxed{\text{Finite set of points}} \rightarrow \boxed{\text{Union of balls}} \rightarrow \boxed{\text{Simplicial complex}}$

\+ The topology of a compact set K can be computed from the Cech complex of a sample P of K

– The Cech complex is huge ($O(n^d)$) and very difficult to compute

– The Cech complex is in general not homeomorphic to K (a triangulation of K)

– The Cech complex cannot be realized in general in the same space as K

∼ Replace the α-Cech complex by the α-complex (less big and embedded)

Looking for small and faithful simplicial complexes

Need to compromise

- Size of the complex
  - can we capture the intrinsic dimensionality ?
- Efficiency of the construction algorithms and of the representations
  - can we avoid the exponential dependence on d ?
  - can we minimize the number of simplices ?
- Quality of the approximation
  - Homotopy type & homology (RIPS complex, persistence)
  - Homeomorphism (Delaunay-type complexes)

Submanifolds of $\mathbb{R}^d$

A compact subset $\mathbb{M} \subset \mathbb{R}^d$ is a submanifold without boundary of intrinsic dimension $k < d$, if any $p \in \mathbb{M}$ has an open (topological) k-ball as a neighborhood in $\mathbb{M}$

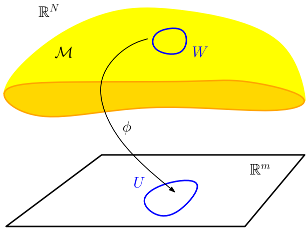

A curve a 1-dimensional submanifold
A surface is a 2-dimensional submanifold

Voronoi diagram and Delaunay complex

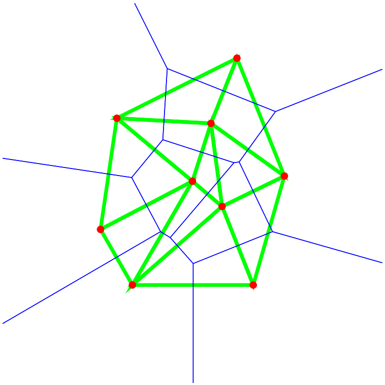

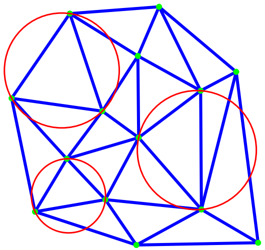

Delaunay complex : $\quad \mathrm{Del(P)} =$ nerve of $\mathrm{Vor(P)}$

Equivalently, $\mathrm{Del(P)}$ is the collection of simplices with an empty circumscribing ball

The curses of Delaunay triangulations in higher dimensions

- Restricted to Euclidean space (see otherwise Mael's talk)
- Computing DT is restricted to low dimensions

  (The number of simplices grows exponentially with d even if the vertices lie on a curve !)
- 3 and higher dimensional Delaunay triangulations are not thick even if the vertices are well-spaced

3D Delaunay Triangulations are not thick even if the vertices are well-spaced

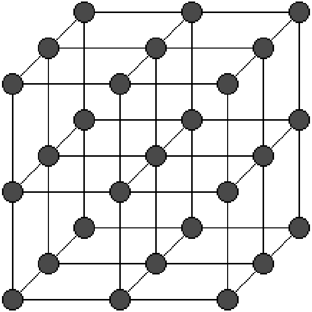

- Each square face can be circumscribed by an empty sphere
- This remains true if the grid points are slightly perturbed therefore creating thin simplices

Badly-shaped simplices

Badly-shaped simplices lead to bad geometric approximations

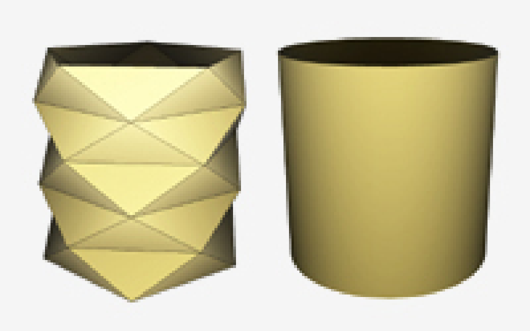

Bad consequences in rendering, numerical simulations, volume calculation and more...

see also [Cairns], [Whitehead], [Munkres], [Whitney]

Simplex quality

Altitudes

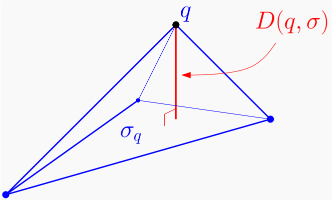

If σ_q, the face opposite q in σ is protected, The *altitude* of q in σ is

$$D(q, \sigma) = d(q, \text{aff}(\sigma_q)),$$

where σ_q is the face opposite q.

Definition (Thickness [Cairns, Whitney, Whitehead et al.])

The *thickness* of a j-simplex σ with diameter $\Delta(\sigma)$ is

$$\Theta(\sigma) = \begin{cases} 1 & \text{if } j = 0 \\ \min_{p \in \sigma} \frac{D(p,\sigma)}{j\Delta(\sigma)} & \text{otherwise.} \end{cases}$$

Tangent space approximation

Lemma

[Whitney 1957]

If σ is a j-simplex whose vertices all lie within a distance h from a hyperplane $H \subset \mathbb{R}^d$, then

$$\sin \angle(\text{aff}\,(\sigma), H) \leq \frac{2j\,h}{D(\sigma)} = \frac{2h}{\Theta(\sigma)\,\Delta(\sigma)}$$

Corollary

If σ is a j-simplex, $j \leq k$, vert $(\sigma) \subset \mathbb{M}$, $\Delta(\sigma) \leq 2\varepsilon\,\text{rch}(\mathbb{M})$

$$\forall p \in \sigma, \quad \sin \angle(\text{aff}(\sigma), T_p) \leq \frac{2\varepsilon}{\Theta(\sigma)}$$

($h \leq \frac{\Delta(\sigma)^2}{2\,\text{rch}(\mathbb{M})}$ by the Chord Lemma)

Chord lemma

Let x and y be two points of $\mathbb{M}$. We have

1. $\sin \angle(xy, T_x) \leq \frac{\|x-y\|}{2\,\mathrm{rch}(\mathbb{M})}$;
2. the distance from y to T_x is at most $\frac{\|x-y\|^2}{2\,\mathrm{rch}(\mathbb{M})}$.

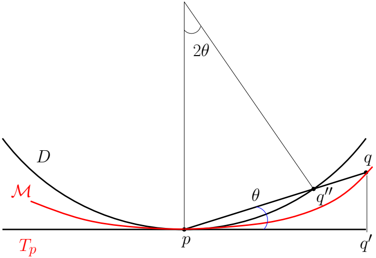

$$\begin{aligned}\|x-y\| &\geq \|x-y''\| \\ &= 2\,\mathrm{rch}(\mathbb{M})\ \sin\angle(xy, T_x)\end{aligned}$$

$$\begin{aligned}\|y-y'\| &= \|x-y\|\ \sin\angle(xy, T_x) \\ &\leq \frac{\|x-y\|^2}{2\,\mathrm{rch}(\mathbb{M})}\end{aligned}$$

The curses of Delaunay triangulations in higher dimensions

- Restricted to Euclidean space (see otherwise Mael's talk)

  $\Rightarrow$ Define local Euclidean triangulations

- Computing DT is restricted to low dimensions

  (The number of simplices grows exponentially with d even if the vertices lie on a curve !)

  $\Rightarrow$ Exploit the fact that $\mathbb{M}$ has an intrinsic dimension $k \ll d$?

- 3 and higher dimensional Delaunay triangulations are not thick even if the vertices are well-spaced

  $\Rightarrow$ Remove flat simplices

Towards Delaunay triangulation of manifolds

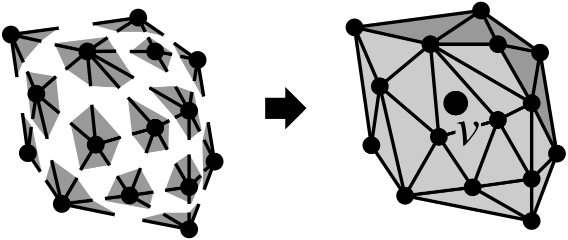

1. Construct local Delaunay triangulations
2. Ensure that the triangulations are stable under small perturbation

   $\Rightarrow$ a simplex belongs to the stars of all its vertices
3. Glue all local triangulations into a single triangulated manifold

The assumptions

- $\mathbb{M}$ is a differentiable submanifold of positive reach of $\mathbb{R}^d$
- The dimension k of $\mathbb{M}$ is small
- $\mathcal{P}$ is an ε-net of $\mathbb{M}$ for a small enough ε
- We assume that we know the tangent space T_p at each $p \in \mathcal{P}$

Local triangulation : $\mathrm{Del}_{T_p}(\mathcal{P})$

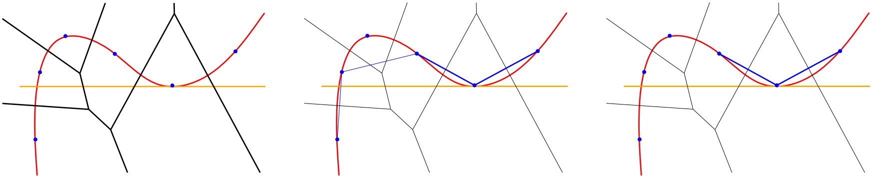

Constructing $\mathrm{Del}_{T_p}(\mathcal{P})$

Given a d-flat $H \subset \mathbb{R}$, $\mathrm{Vor}(\mathcal{P}) \cap H$ is a weighted Voronoi diagram in H

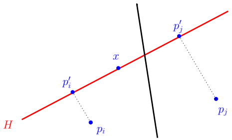

$$\|x - p_i\|^2 \leq \|x - p_j\|^2$$
$$\Leftrightarrow \quad \|x - p'_i\|^2 - \|p_i - p'_i\|^2 \leq \|x - p'_i\|^2 - \|p_j - p'_j\|^2$$

Corollary: construction of Del_{T_p}

$$\psi_p(p_i) = (p'_i, -\|p_i - p'_i\|^2) \qquad \text{(weighted point)}$$

1. project $\mathcal{P}$ onto T_p which requires $O(Dn)$ time
2. construct $\mathrm{star}(\psi_p(p_i))$ in $\mathrm{Del}(\psi_p(p_i)) \subset T_{p_i}$
3. $\mathrm{star}(p_i) \approx \mathrm{star}(\psi_p(p_i))$ (isomorphic)

The tangential Delaunay complex

[Freedman 2002], [B.& Flottoto 2004], [B. Ghosh 2014]

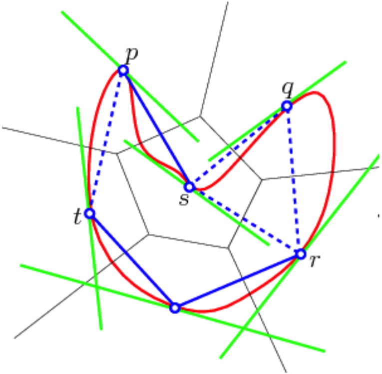

1. Construct the star of $p \in \mathcal{P}$ in the Delaunay triangulation $\mathrm{Del}_{T_p}(\mathcal{P})$ of $\mathcal{P}$ restricted to T_p
2. $\mathrm{Del}_{T\mathbb{M}}(\mathcal{P}) = \bigcup_{p \in \mathcal{P}} \mathrm{star}(p)$

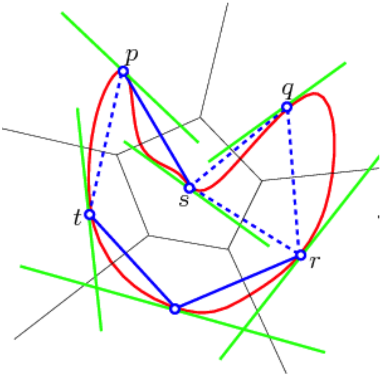

- $+$ $\mathrm{Del}_{T\mathbb{M}}(\mathcal{P}) \subset \mathrm{Del}(\mathcal{P})$
- $+$ $\mathrm{star}(p)$, $\mathrm{Del}_{T_p}(\mathcal{P})$ and therefore $\mathrm{Del}_{T\mathbb{M}}(\mathcal{P})$ can be computed without computing $\mathrm{Del}(\mathcal{P})$
- – $\mathrm{Del}_{T\mathbb{M}}(\mathcal{P})$ is not necessarily a triangulated manifold

Inconsistent configurations

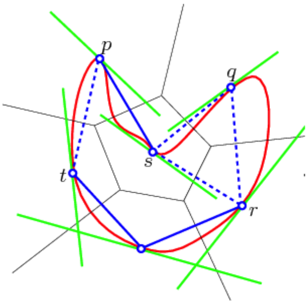

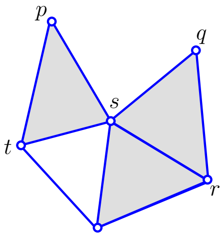

Definition

$\phi = [p_1, p_2, \ldots, p_{k+2}]$

$\tau = \phi \setminus \{p_l\}$

ϕ is an inconsistent configuration witnessed by $p_i, p_j, p_l \in \phi$ if

- $\tau \in \mathrm{star} p_i \quad \Leftrightarrow T_{p_i} \cap \mathrm{Vor}(\tau) \neq \emptyset$
- $\tau \notin \mathrm{star} p_j \quad \Leftrightarrow T_{p_j} \cap \mathrm{Vor}(\tau) = \emptyset$
- $\mathrm{Vor}(p_l)$ is the first Voronoi cell whose interior is hit by $[c_{p_i}(\tau) \to c_{p_j}(\tau)]$

An IC is a $(k+1)$-simplex of $\mathrm{Del}(\mathbb{M})$

Inconsistent configurations are not thick

1. $\tau \in \text{star}(p_i) \Rightarrow B(c_{p_i}(\tau) \cap \mathcal{P} = \emptyset$
2. $\tau \notin \text{star}(p_j) \Rightarrow B(c_{p_i}(\tau) \cap \mathcal{P} \ni p$

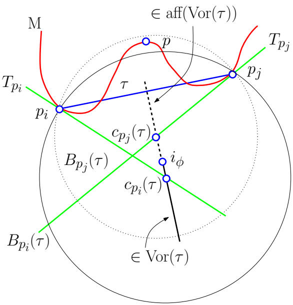

if τ is small and thick

$\Rightarrow c_i$ and c_j are close & $\text{aff}(\tau) \approx T_{p_i} \approx T_{p_j}$

$\Rightarrow \phi := \tau * p$ is not thick

Reconstruction of smooth submanifolds

1. For each vertex v, compute the star $\mathrm{star}(p)$ of p in $\mathrm{Del}_p(\mathcal{P})$
2. Remove inconsistencies among the stars by weighting the points
3. Glue the stars to obtain a triangulation of $\mathcal{P}$

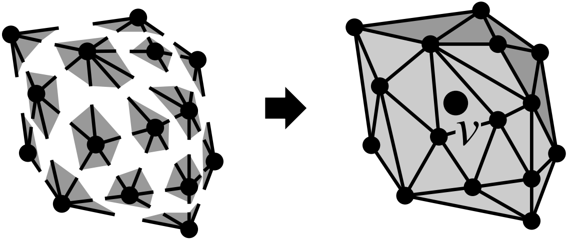

Manifold reconstruction algorithm

Algorithm 1 Manifold_reconstruction(P $= \{p_0, \ldots, p_n\}, \eta_0)$

// Initialization
for $i = 1$ to n **do**
  calculate the local neighborhood $LN(p_i)$
for $i = 1$ to n **do**
  $\omega(p_i) \leftarrow 0$
Build the full unweighted complex $K^\omega(\mathsf{P}) = \mathrm{Del}^\omega_{T\mathbb{M}}(\mathcal{P}) \bigcup IC$

// Weight assignment to remove inconsistencies
for $i = 1$ **to** n **do**
  $\omega(p_i) \leftarrow \mathbf{weight}(p_i, \omega)$
  update $K^\omega(\mathsf{P})$ // (locally in $LN(p_i)$

output : $\hat{\mathbb{M}} \leftarrow \mathrm{Del}^\omega_{T\mathbb{M}}(\mathsf{P})$

Hypotheses

- $\mathbb{M}$ is a differentiable submanifold of positive reach of dim. $k \subset \mathbb{R}^d$
- $\mathcal{P}$ is an ε-net of $\mathbb{M}$ for a small enough ε

Theorem

Under the Hypotheses, the algorithm terminates and $\hat{\mathbb{M}}$ contains no inconsistent configurations

$\hat{\mathbb{M}}$ is a PL simplicial k-manifold

Lemma

Let P be an ε-sample of a manifold $\mathbb{M}$ and let $p \in$ P. The link of any vertex p in $\hat{\mathbb{M}}$ is a topological $(k-1)$-sphere

Proof :

1. Since $\hat{\mathbb{M}}$ contains no inconsistencies, the star of any vertex p in $\hat{\mathbb{M}}$ is identical to $\mathrm{star} p$, the star of p in $\mathrm{Del}_p(\mathsf{P})$

2. $\mathrm{Del}_p(\mathsf{P}) \subset \mathbb{R}^d \; \approx \; \mathrm{Del}(\psi_p(\mathsf{P})) \subset T_p \quad \Rightarrow \quad \mathrm{star} p \approx \mathrm{star}_p(p)$

3. $\mathrm{star}_p(p)$ is a k-dimensional triangulated topological ball (general position)

4. p cannot belong to the boundary of $\mathrm{star}_p(p)$

(the Voronoi cell of $p = \psi_p(p)$ in $\mathrm{Vor}(\psi_p(\mathsf{P}))$ is bounded)

Guarantees

Approximation guarantees

- $\hat{\mathbb{M}}$ is a PL simplicial k-manifold
- $\hat{\mathbb{M}} \subset$ tub$(\mathbb{M}, O(\varepsilon^2)\mathrm{rch}(\mathbb{M}))$
- The angles between the facets and the tangent spaces of $\mathbb{M}$ are $O(\varepsilon)$
- $\hat{\mathbb{M}}$ is homeomorphic to $\mathbb{M}$

Complexity of the algorithm

- No d-dimensional data structure $\Rightarrow$ linear in d
- exponential in k

Reconstructing a Riemannian surface in $\mathbb{R}^8$

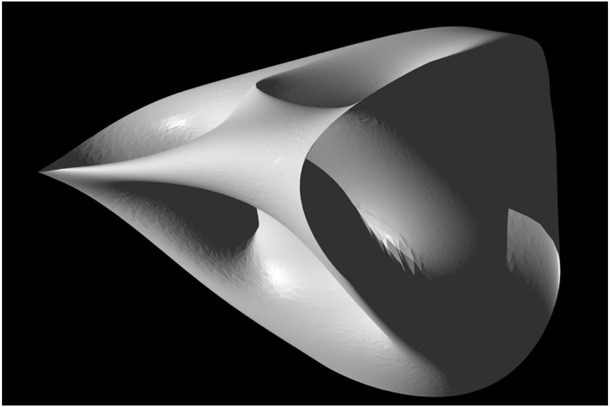

Definition

The restricted Delaunay triangulation $\mathrm{Del}_X(\mathcal{P})$ to $X \subset \mathbb{R}^d$ is the nerve of $\mathrm{Vor}(\mathsf{P}) \cap X$

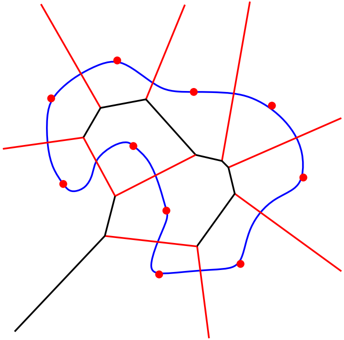

If P is an ε-sample, any ball centered on X that circumscribes a facet f of $\mathrm{Del}_X(\mathcal{P})$ has a radius $\leq \varepsilon\, \mathrm{lfs}(c_f)$

Restricted Delaunay triangulation [Chew 93]

Definition

The restricted Delaunay triangulation $\mathrm{Del}_X(\mathcal{P})$ to $X \subset \mathbb{R}^d$ is the nerve of $\mathrm{Vor}(\mathsf{P}) \cap X$

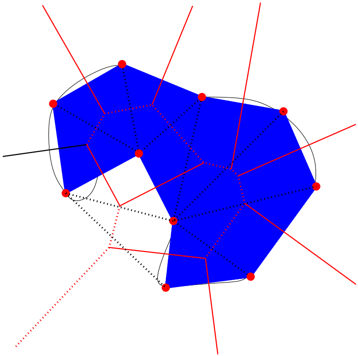

Delaunay triang. restricted to surfaces

[Amenta et al. 1998-], [B. & Oudot 2005]

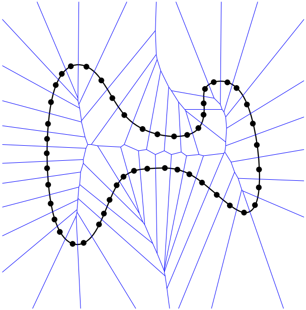

If

- $\mathcal{S} \subset \mathbb{R}^3$ is a compact surface of positive reach without boundary
- $\mathcal{P}$ is an ε-net, ε small enough

then

- $\mathrm{Del}_{|S}(\mathcal{S})$ provides good estimates of normals
- There exists a homeomorphism $\phi : \mathrm{Del}_{|S}(\mathcal{P}) \rightarrow \mathcal{S}$
- $\sup_x(\|\phi(x) - x\|) = O(\varepsilon^2)$

Thickness and angle bounds are automatic for *triangles*

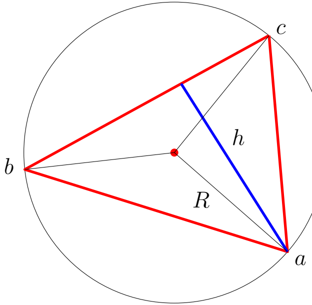

$$h = ab \ \times \ sinb = \frac{ab \times ac}{2R}$$

$$\Theta(abc) = \frac{h}{2\Delta} \geq \frac{h}{4R} \geq \frac{\phi_0^2}{32R} \geq \frac{{\phi_0}^2}{32\varepsilon \, \mathrm{rch}(\mathcal{S})}$$

$$\sin b \geq \frac{\phi_0}{2\varepsilon \, \mathrm{rch}(\mathcal{S})}$$

Surface mesh generation by Delaunay refinement

[Chew 1993, B. & Oudot 2003]

$\phi : S \to \mathbb{R}$ = Lipschitz function
$\forall x \in S,\ 0 < \phi_0 = \bar{\eta}_0\, \varepsilon\, \mathrm{rch}(\mathcal{S}) \leq \phi(x) < \varepsilon\, \mathrm{lfs}(x)$

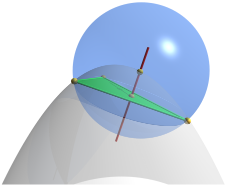

ORACLE : For a facet f of $\mathrm{Del}_{|S}(\mathcal{P})$,
return c_f, r_f and $\phi(c_f)$

A facet f is bad if $r_f > \phi(c_f)$

Algorithm

INIT compute an initial (small) sample $\mathcal{P}_0 \subset S$

REPEAT IF f is a bad facet
insert_in_Del3D(c_f) ,
update $\mathcal{P}$ and $\mathrm{Del}_{|S}(\mathcal{P})$

UNTIL no bad facet remains

The meshing algorithm in action

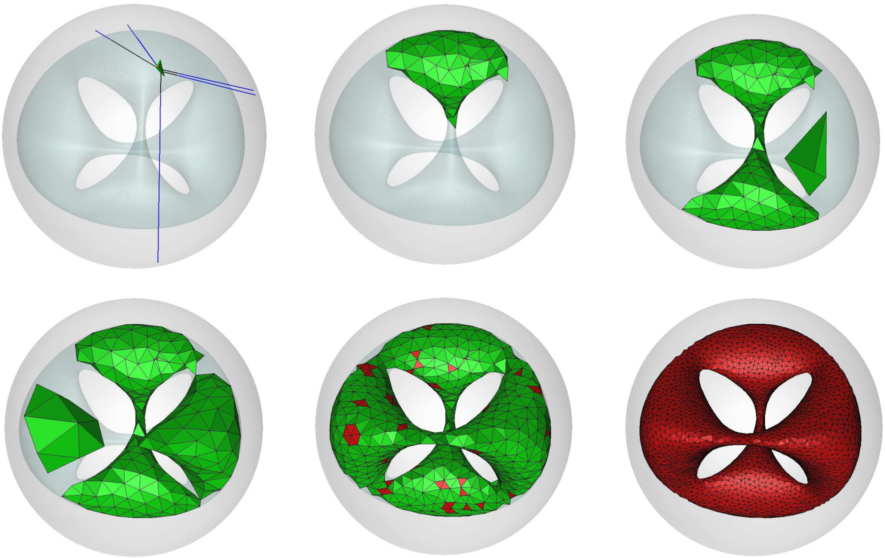

Applications

- Implicit surfaces $f(x, y, z) = 0$
- Isosurfaces in a 3d image (Medical images)
- Triangulated surfaces (Remeshing)
- Point sets (Surface reconstruction)

see cgal.org, CGALmesh project

Results on smooth implicit surfaces

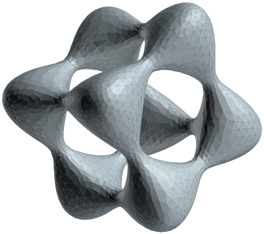

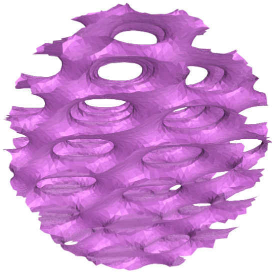

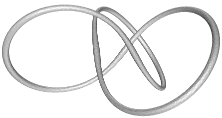

Meshing 3D domains

Input from segmented 3D medical images

[INSERM]

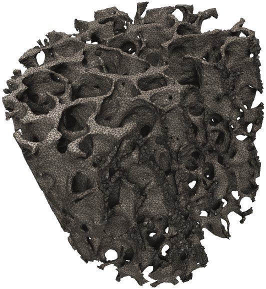

[SIEMENS]

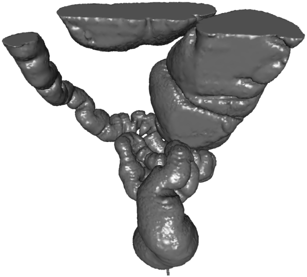

Comparison with the Marching Cube algorithm

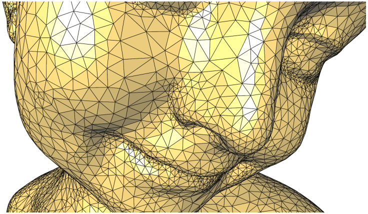

Delaunay refinement

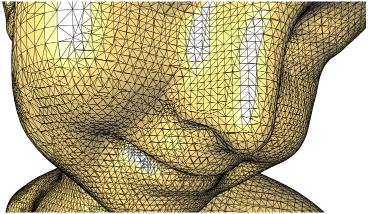

Marching cube

Meshing with sharp features

A polyhedral example

.

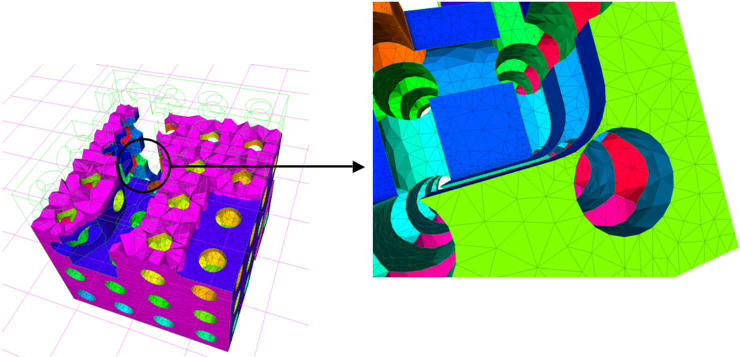

Meshing 3D multi-domains

Input from segmented 3D medical images [IRCAD]

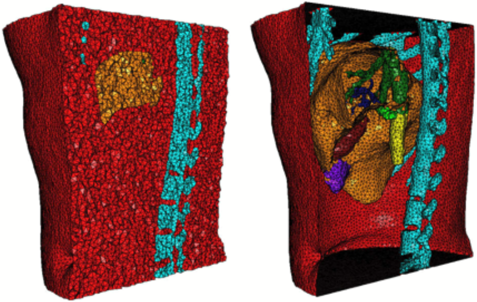

Size bound (mm)	vertices nb	facets nb	tetrahedra nb	CPU Time (s)
16	3,743	3,735	19,886	0.880
8	27,459	19,109	159,120	6.97
4	199,328	76,341	1,209,720	54.1
2	1,533,660	311,420	9,542,295	431

Surface reconstruction from unorganized point sets

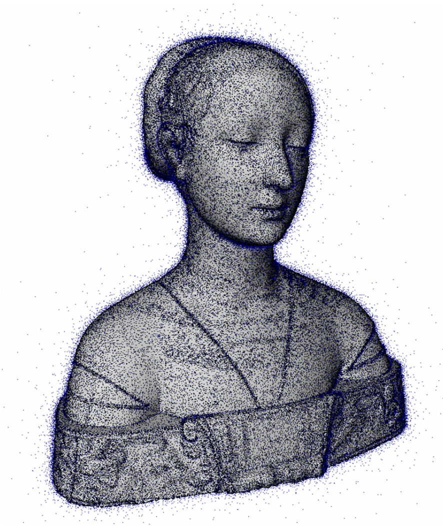

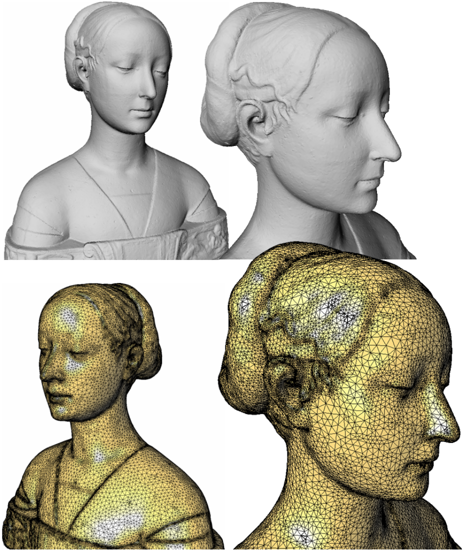

Courtesy of P. Alliez